

SECTION 4: BINARY DATA SECTION (BDS).

The BDS contains the packed data and the binary scaling information needed to reconstruct the original data from the packed data. The required decimal scale factor is found in the PDS, above. The data stream is zero filled to an even number of octets.

Octet no.	
1 - 3	Length in octets of binary data section
4	Bits 1 through 4: Flag - See Table 11 Bits 5 through 8: Number of unused bits at end of Section 4.
5 - 6	The binary scale factor (E). A negative value is indicated by setting the high order bit (bit No. 1) in octet 5 to 1 (on).
7 - 10	Reference value (minimum value); floating point representation of the number.
11	Number of bits into which a datum point is packed
12 -nnn	Variable, depending on octet 4; zero filled to an even number of octets.
14	Optionally, may contain an extension of the flags in octet 4. See Table 11.

Here are some of the various forms the binary data can take; the flag table in BDS octet 4, possibly extended into octet 14, identifies which variant is in use.

Grid-point data - Simple packing

Here the data simply begin in octet 12 and continue, packed according to the simple packing algorithm described above, without any particular regard for computer "word" boundaries, until there is no more data. There may be some "zero-fill" bits at the end.

If all the data in a grid point field happen to have the same value, then all of the deviations from the reference value are set to zero. Since a zero value requires no bits for packing, octet 11 is set to zero, thus indicating a field of constant data, the value of which is given by the reference value. Under these circumstances, octet 12 is set to zero (the required "zero fill to an even number of octets") and bits 5-8 of octet 4 contain an 8. The number of data points in the field is implied by the grid identification given in the PDS and/or the GDS and BMS.

Spherical Harmonic Coefficients - Simple packing

Octets 12-15 contain the real part of the (0,0) coefficient in the same floating point format as the reference value in octets 7-10. The imaginary part of the (0,0) coefficient, mathematically, is always equal zero. Octets 16 to the end contain the remaining coefficients packed up as binary data with the same sort of scaling, reference value, and the like, as with grid-point numbers. Excluding the (0,0) coefficient, which is usually much larger than the others, from the packing operation means that the remaining coefficients can be packed to a given precision more efficiently (fewer bits per word) than would be the case otherwise.

Grid-Point Data - Second Order or Complex Packing

Before laying out where the various second order values, sub-parameters, counters, and what have you, go, it is appropriate to describe the second order packing method in an algorithmic manner.

Referring back to the description of simple packing, the encoding method is the same up to part way through the fourth step, stopping just short of the actual packing of the scaled integers into the "words" of either a pre-specified or calculated bit length.

The basic outline of second order packing is to scan through the array of integers (one per grid point, or possibly less than that if the Bit Map Section has been employed to discard some of the null value points) and seek out sub-sections exhibiting relatively low variability within the sub-section. One then finds the (local) minimum value in that sub-section and subtracts it from the ("first order") integers in that sub-section, which leave a set of "second order" integers. These numbers are then scanned to find the maximum value, which in turn is used to specify the minimum bit width for a "word" necessary to contain the sub-section set of second order numbers.

The term "first order" in this context refers to the integer variables that result from subtracting the overall (global) minimum from the original variables and then doing all scaling and rounding; "second order" refers to the variables that result from subtracting the local minimum from the sub-set of first order variables. No further scaling is necessary or appropriate.

The sub-section set of numbers are then packed into "words" of the just determined bit length. The overall savings in space comes about because the second order values are, usually, smaller than their first order counterparts. They have, after all, had two minima subtracted from the original values, the overall minimum and the local minimum, where the first order values have had only the overall minimum subtracted out. There is no guarantee, however, that the second order packing will compress a given field to a greater degree than the first order packing. If the first order field of integers is highly variable, or generally close to zero, then there will be no gain in compression. But if the field shows long runs of small variation, particularly if some of the runs are constant (zero variability), then the second order packing will contribute to the compression.

The process then repeats and a whole collection of sub-sections are found, their local minima are subtracted, etc. One of the tricky parts of this process is defining just what is meant by a "sub-section of low variability". The WMO Manual is silent on this as it only describes how the sub-sections and their ancillary data are to be packed in the message. The U.S. National Weather Service, the U.K. Meteorological Office, the European Centre for Medium-Range Weather Forecasts, and probably other groups have, independently, designed selection criteria and built them into GRIB encoders. It is beyond the scope of this document to attempt to describe them in any detail. These groups have all expressed their willingness to share their GRIB encoders with any who ask for them.

Before laying out where the second order values, etc., are placed in a message, we had best review just what information has to be saved. We need to include the following information:

- 1) How many sub-sections there are;
- 2) Where does each sub-section begin;
- 3) Where does each sub-section end; or, how many data points are in each sub-section;
- 4) What is the local minimum value (a first order value) that was found for each sub-section;
- 5) What is the bit width of the collection of first order values (the local minima) found for each sub-section;
- 6) What are the second order values for each sub-section;
- 7) What are the bit widths of the second order values appropriate for all the sub-section; and, finally,
- 8) Sufficient information to specify where the above information is located.

A moments consideration (a long moment, perhaps) will satisfy the reader that the information given will be sufficient to reconstruct the original data field.

The information needed for points 2) and 3), the beginning and end of the sub-sections, is presented in the form of a bit map, called a "secondary bit map" to distinguish it from the bit map (optionally) contained in the BMS. There is one bit for each grid point containing data, ordered in the same way as the grid is laid out. The "primary" bit map, the BMS bit map, may have been used to eliminate data at points where the data are meaningless - only the remaining "real" data points are matched by the bits in the secondary bit map. This possibility is understood to exist throughout the following discussion. The start of each sub-section is indicated by the corresponding bit set to "on" or to a value of 1. Clearly, the first bit in the secondary bit map will always be set on, since the first data point must be the start of the first sub-section. (If it is not, then something is wrong somewhere. Unfortunately it is not always easy to tell just where the error occurred.) The secondary bit map is then no more than a collection of 1s and 0s, indicating the start and the extent of each sub-section. It would be possible to scan through the secondary bit map and determine how many sub-sections there are;

however, this number is explicitly included in the GRIB message to save one the trouble, and to serve as an internal self-checking mechanism.

At long last, then, here is the layout of the information, with further explanatory notes, when second order packing has been employed:

Octet no.	Content
1-3	Length in octets of binary data section
4	Bits 1 through 4: Flag - See Table 11 Bits 5 through 8: Number of unused bits at end of Section 4.
5-6	The binary scale factor (E). A negative value is indicated by setting the high order bit(bit No. 1) in octet 5 to 1 (on).
7-10	Reference value (minimum value); floating point representation of the number. This is the overall or "global" minimum that has been subtracted from all the values.
11	Number of bits into which a datum point is packed. This width now refers to the collection of first order packed values that serve as the local minimum values, one for each sub-section. It is determined in the same manner as for the simple (first order) packing.
12-13	N1 - Octet number, relative to the start of the BDS, at which the collection of first order packed numbers begins, i.e. the collection of local minimum values.
14	The flags that are an extension of octet 4. See Code Table 11.
15-16	N2 - Octet number, relative to the start of the BDS, at which the collection of second order packed numbers begins.
17-18	P1 - The number of first order packed values, the local minima. This number is the same as the number of sub-sections.
19-20	P2 - The number of second order packed values actually in the message. This is the number of data points as (possibly) modified by the bit map in the BMS, if any, and/or reduced by the number of identical points collapsed together by the run-length encoding (see below).
21	Reserved

22-(xx-1)	Width(s), in bits, of the second order packed values; each width value is contained in 1 octet. There are as many width values, or octets, as there are sub-sections, P1 of them. However, there may be but one such value under special circumstances; see below. Also, the width value for a particular sub-section may perfectly well be zero.
xx-(N1-1)	Secondary bit map, one bit for each data point. It will be P2 bits long, then padded to an even number of octets with binary 0.
N1-(N2-1)	P1 first order packed values, the local minima, each held in a "word" of bit-length found in octet 11, then padded to a whole number of octets with 0s.
N2-...	P2 second order packed values. There is no "marking" of the sub-sections here; all the sub-section second order values are placed in a continuous string of bits. The bit-length of the "words" holding the values will change from place to place but again this has to be determined by reference to the other information.

As usual, there may be padding by binary 0 bits sufficient to bring the entire section to an even number of octets.

There are a small number of special cases and variations on the above layout:

If the bit-width for a sub-section is zero, then no second order values for that sub-section are included in the part of the message starting at octet N2. The value of P2 will reflect the absence of those points. This will happen if all the first order values in the sub-section are identical. This is a form of "run-length encoding" and contributes greatly to packing efficiency if the original data contains strings of constant value (including zero).

Under some circumstances, it may turn out that there is no need to use different bit-widths for each of the sub-sections. In that case, a flag is set in bit 8 of the extended flags found in octet 14 (see table 11) indicating that all the sub-sections are packed with the same bit-width, and that the single value will be found in octet 22.

Row by row packing is defined as selecting entire rows (or columns) to serve as sub-sections, without regard to "variability" determinations. It can have some compression value. If row by row packing is employed, this is indicated by setting a flag in bit 7 of the extended flags found in octet 14 (see table 11) and NOT including the secondary bit map in the message. It is unnecessary since the length of the rows (columns) is known from the grid specifications given elsewhere in the message.

TABLES FOR THE BDS

TABLE 11. FLAG
(BDS Octet 4 and, optionally, 14)

Bit	Value	Meaning
1	0	Grid point data
	1	Spherical Harmonic Coefficients
2	0	Simple packing
	1	Second order ("Complex") Packing
3	0	Original data were floating point values
	1	Original data were integer values
4	0	No additional flags at octet 14
	1	Octet 14 contains flag bits 5 - 12

The following gives the meaning of the bits in octet 14 ONLY if bit 4 is set to 1. Otherwise octet 14 contains regular binary data.

5		Reserved (set to 0)
6	0	Single datum at each grid point
	1	Matrix of values at each grid point
7	0	No secondary bit maps
	1	Secondary bit maps present
8	0	Second order values have constant width
	1	Second order values have different widths
9-12		Reserved (set to 0)

Notes:

- (1) Bit 3 is set to 1 to indicate that the original data were integers; when this is the case any non-zero reference values should be rounded to an integer value prior to placing in the GRIB BDS.
- (2) Bit 4 is set to 1 to indicate that bits 5 to 12 are contained in octet 14 of the data section.
- (3) Although GRIB is not capable of representing a matrix of data values at each grid point, the meaning of bit 6 is retained in anticipation of a future capability.
- (4) When secondary bit maps are present in the data (used in association with second order packing) this is indicated by setting bit 7 to 1.

- (5) When octet 14 contains the extended flag information octets 12 and 13 will also contain "special" information; the actual data will begin in a subsequent octet. See above.

At present, the "extension" of Table 11 into octet 14 and the associated "advanced" features of GRIB are limited to spherical harmonics and second order("complex") packing in the National Weather Service. Additional variations are included in the WMO Documentation.